

Amendments to the drawings,

Proposed amendments to the Drawings are shown in the attached Replacement Sheets (Figs. 3 and 6A-D). Formal drawings of the Replacement Sheets will be submitted via mail, upon indication from the Examiner that the proposed amendments are acceptable.

Remarks

Status of application

Claims and 1-53 are pending in the subject application. The claims stand rejected in view of cited prior art, and as well as certain informalities (discussed below). By this Amendment, the claims have been amended to address the rejections. Re-examination and reconsideration of the claims are respectfully requested.

The invention

In accordance with the present invention, operation of an e-mail system is modified to incorporate a flow control filter (service). During processing of incoming e-mail, each child MTA process (that is created to handle a particular new connection) connects to the flow control filter service, so that it can interact with the service during arrival of a message. This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello" handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail. Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered.

The overall methodology of operation may be summarized as follows. The following method steps occur in the context of an incoming message that is being processed by the e-mail system (i.e., MTA forking has already occurred) and now the system is ready to evoke the services of the flow control filter of the present invention.

Invocation of the flow control filter begins with the MTA (i.e., a child MTA of the original (parent) listener) connecting to the flow control filter (e.g., using Sendmail Milter protocol); the filter accepts the connection. The MTA and the filter perform a handshake sequence, including feature and parameter negotiation. At the conclusion of the handshake sequence, a new thread is created (i.e., in the flow control engine) for processing the new connection/message. Now, the MTA passes to the filter the corresponding connection information (e.g., IP address and host name) of the sending MTA. Based on the connection information, the filter may look up matching class data from the configuration file. In the event that no matching class data is found, the filter will assume unrestricted access for the host and therefore will accept the connection and message. In that case, the flow control engine thread handling the connection may terminate, as there is no further filtering work to be done for this incoming connection and message; the MTA proceeds normally with no further interaction with the filter. Otherwise, the method proceeds to the following filtering steps. The method tests whether class limits have been reached. In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current connection count. Otherwise (i.e., false), the method terminates with the filter rejecting the connection and returning an administrator-defined error code. In the event that the process did not terminate, the MTA reports the sender information to the filter; this occurs in response to the MAIL FROM SMTP phase.

The method notes the sender (i.e., who is the sender) in the class. The administrator-defined class may include, for example, a sender-based parameter indicating that the filter should note the number of unique senders that have arrived in a given timeframe for this particular host (of the class). In a manner similar to above, the method tests whether class' sender limits have been reached. In the event that limits have not been reached (i.e., true), the filter instructs the MTA to continue and increments the current unique sender totals. Otherwise, the method terminates with the filter rejecting the message (returning any administrator-defined error code). In the event that the filtering process has not terminated based on sender information, the method proceeds to test recipient (RCPT TO) information. The configuration file allows the administrator to

define a class that limits the number of unique recipients received for that class, over any given time span. As a given message may have multiple recipients, the step repeats for each recipient (information) of the message. As before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). Otherwise, the method updates the totals and proceeds.

The MTA reports the message body, which may be transmitted as one or more blocks. The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. The MTA reports end of message for the current incoming message. The method compares the message size against class limits specified in the configuration file. Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code). Otherwise, the incoming message has passed all filters and is accepted. Now, the method may repeat for other incoming messages.

This approach may be easily scaled, for application on a site-wide basis. In that instance, the flow control filter service monitors the children processes for a number of e-mail servers at a given site. In such a configuration, the flow control filter service would apply policy on a global (site) basis, instead of on a per server basis.

General

A. Information Disclosure Statement (IDS)

The Examiner has complained that, "The listing of references in the specification is not a proper information disclosure statement. 37 CFR 1.98(b) requires a list of all patents, publications, or other information submitted for consideration by the Office." The list of references (RFC's -- Request for Comments) in the specification was never intended to serve as an IDS. The list simply indicates the general background of the invention and the general state of the e-mail art. In any event, the references have been submitted for the convenience of the Examiner in an IDS filed via facsimile transmission on March 7, 2005. A downloadable PDF copy of the filed IDS is also made available for the Examiner's convenience at:

<http://www.Smart-IPLaw.com/fax/SMI0005.00.IDS02.Fax.03.07.2005.pdf>

B. Drawings

In a first objection to the Drawings, the Examiner complains that item 301 is not show in Fig. 3. Actually, "301" is intended to refer collectively to subitems 301a, 301b, 301c, and 301d. A revised Fig. 3 which shows this grouping explicitly with "301" is submitted herewith.

In a second objection to the Drawings, the Examiner notes that item 301d in Fig. 3 is not mentioned in the specification. This omission has been corrected by amending the specification to change

application software or "programs" 301 (e.g., 301a, 301b, 301c)

to:

application software or "programs" 301 (e.g., 301a, 301b, 301c, 301d)

In a third objection to the Drawings, the Examiner complains that, "The drawings are objected to because Figure 6a-6d are presented in the form of a flowchart. Some IF/ELSE conditions exist in the drawings that are not treated as such in a normal flowchart." Figs. 6A-6D have been revised to present the method as format-neutral method steps or pseudocode, thus eliminating the need to illustrate the method steps as a flowchart having a particular format. The Brief Description text of Figs. 6A-6D has likewise been amended to reflect the change.

C. Specification

The Examiner objects to the specification because it contains an embedded hyperlink and/or other form of browser-executable code. Specifically, "http://" has been removed from the paragraph beginning at page 5, line 20, the paragraph beginning at page 8, line 15, and the paragraph beginning at page 17, line 2. Additionally, "ftp://" is removed from the paragraph beginning at page 8, line 29. Accordingly, embedded hyperlink and/or other form of browser-executable code is now removed.

D. Claims Rejections, 35 U.S.C. 112, second paragraph

1. Claims 2, 3, 5, 8, 11, 14, 15, 29, 42-46, and 53

The Examiner has objected to the use of "a period of time" in claims 2, 3, 5, 8, 11, 14, 15, 29, 42-46, and 53. The Examiner states that:

A period of time is an indefinite limitation. It would cause one of ordinary skill in the networking art undue experimentation to implement the invention because of the unknown quantity of "a period of time". For purposes of compact prosecution, "a period of time" will be treated as occurring over any period of time.

As the Examiner correctly surmises, the phrase "a period of time" as used herein refers to any period of time as specified in the configurable policy rules, which is what the claim language specifically requires. Applicant's claim 2, for example, sets forth the limitation:

2. The method of claim 1, wherein said configurable policy rules specify a maximum number of connections permitted by a given domain over a period of time.

As set forth above, the period of time is established as part of the configurable policy rules, which have been authored by the user (administrator) using the system of the present invention. It refers to some period, interval, or span of time, such as a period of 1 hour, 30 minutes, 15 hours, five days, one week, 5 years, or whatever arbitrary time period one wishes to specify in the configurable policy rules. The point of the claim is a maximum number of connections occurs over some given period of time -- however long or short it is, is left to the discretion of the user (administrator) configuring the configurable policy rules.

In an effort to address the Examiner's objection, the claims have been amended to clarify that the period of time is the "desired period of time" that is specified in the configurable policy rules. The foregoing amendment is not offered for purposes of distinguishing over the prior art, but instead is merely offered to clarify text that the

Examiner cites as indefinite. If the Examiner finds that the amendment does not address his concern, it is respectfully requested that he further clarify the nature of the complained-about indefiniteness, particularly in light of the USPTO's own patent database showing that the phrase "a period of time" appears in the claims of several tens of thousands of U.S. patents issued since 1976.

2. Claims 10, 13, and 37

Claims 10, 13, and 37 contain informalities in the form of typographical errors and incorrect antecedents referencing, as noted by the Examiner. The claims have been amended to correct these informalities. Specifically, in claim 10, "said sender information" should have been "said recipient information". In claim 13, "said sender information" should have been "said e-mail message body data". In claim 37, "RCPT FROM" should have been "RCPT TO".

Prior art rejections

A. Section 102: Srivastava et al.

Claims 1, 16-17, 21-22, 25, 27-28, 31-34, 41, 47-48 and 50-53 stand rejected under 35 U.S.C. 102(e) as being anticipated by Srivastava et al. (U.S. Patent No. 6,374,292), hereinafter referred to "Srivastava"). The Examiner's rejection of claim 1 is representative:

Regarding claim 1, Srivastava discloses a method for processing an incoming e-mail message that is being received from another domain, the method comprising: receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system; in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis; comparing the request from the particular domain against configurable policy rules; and denying the request if any of said policy rules would be violated. [Srivastava discloses using a process to define a particular domain in an email server. An individual domain can be configured to allow all mail to be received if the state of the domain is active (establish a new connection) or if the state of the domain is inactive the domain is suspended from routing mail (denying the request). See Srivastava, column 7, lines 36-59. Srivastava further discusses using a

multithreaded process, with each thread handling a connection. Examiner considers this to be equivalent to creating a second process for handling a new connection. Srivastava further states that using a single multithreaded process is beneficial by maximizing performance and stability and by minimizing system resource usage. See Srivastava, column 5, lines 9-15.] By this rationale claim 1 is rejected.

As shown below, Applicant's claimed invention may be distinguished on a variety of grounds.

Srivastava refers to a package of software running on a mail server which governs which specific e-mail related services are offered to users in a particular domain or set of domains. In fact, all of the services described in Srivastava are at the upper OSI layers, i.e., "Presentation" and "Application". Chief among these is virtual domain hosting, whereby one server is given the ability to send and receive e-mail for a variety of otherwise unrelated domains, and to thereby provide e-mail services for users within those domains. The intent of Srivastava is to give the ISP (Internet Service Provider) the means to delegate these services to the administrators of those respective domains without also maintaining individual e-mail servers for each domain or granting machine-wide administrative access to lots of disparate organizations, which of course are prospects with very serious scaling implications.

Applicant's flow control invention, since it operates below the "Presentation" layer in the OSI model, has little knowledge of which domains are stored on the server(s) it protects. It is instead operating in the "Session" layer, with a small amount of information made available to it from lower layers. Moreover it has no knowledge of any protocol or service related to e-mail other than SMTP. It may even not be running on a server providing SMTP service at all. Instead, it has knowledge of the origin of the TCP connection being made to an SMTP service, and makes use of this knowledge to classify the connection and thereby moderate the flow of e-mail into or out of a system. While Srivastava's invention is more geared toward assistance and delegation of administration of e-mail services, Applicant's flow control invention serves to prevent domination of network and system resources by a particular e-mail source.

The flow control invention is designed to moderate the flow of e-mail inbound.

An example of this might be to limit the maximum number of e-mails, connections over time, or simultaneous connections coming from a given domain (say, e.g., "prodigy.net") so that the impact of a sudden surge in spam or a deliberate denial-of-service attack can be detected and quashed without impacting the flow of e-mail from other domains. An outbound policy can also be applied, so that for example a user's desktop infected with a virus that then tries to send e-mail to hundreds or thousands of users is blocked when it is detected. Both of these concepts are known commonly as "traffic shaping", although that term often describes a function implemented by routers at the OSI "Transport" layer and below. This is in contrast to Srivastava, which is essentially an e-mail server provisioning system, allowing a machine or set of machines to be "sliced up" in such a way that many "virtual" domains can be served by a small number of hosts (or one), and the services provided to those domains can be managed by the domain's respective owners.

Turning now to the Examiner's specific basis for rejection, the Examiner analogizes Applicant's flow control filter to Srivastava's virtual domains (Srivastava, Col. 7, lines 36-59).

Referring now to FIG. 4, showing a flowchart that details a process 500 for defining a virtual domain in accordance with an embodiment of the invention. The process 500 begins at 502 by defining a virtual domain node in the DIT. Once the virtual domain node has been defined, corresponding routing table entries are defined at 504 and at 506, various virtual domain are stored at the virtual domain node. It should be noted that the various virtual domain include a list of services permitted the domain. Such services include IMAP, MAPS, POP3, POP3S, SMTP which in some cases requires presentation of credentials. Other of the services include identification of a domain administrator who is authorized to manage the particular virtual domain which includes setting particular user-level for particular users in the domain. These services also include designation of a virtual domain postmaster who identifies email message delivery problems, and a state of the domain.

As clearly shown above, Srivastava at this point is describing the creation of a virtual domain. The purpose of a "virtual domain" is discussed by Srivastava: "It is therefore desirable that an email service provider be able to offer email services to multiple

organizations each of which has their own virtual domain and to support the ability to define such domains in the directory and host them on a shared mail server." This is not the same as Applicant's limitation of "said second process being connected to a flow control filter providing filtering on a per-domain basis," which is able to block inbound e-mail traffic -- or allow e-mail traffic -- from different domains, depending on whether a particular given domain is complying with the "configurable policy rules."

For example, if a given domain in Srivastava's system is specified to be a virtual domain for which e-mail services are allowed, then Srivastava's system would permit all e-mail traffic from that virtual domain -- regardless of whether some of that traffic is coming from a user machine engaged in spam or a deliberate denial-of-service attack. There certainly is no mention or passing suggestion in Srivastava that his system also includes some sort of adaptive filter that would then somehow further monitor the virtual domains as they make connections to determine whether the connections for e-mail traffic originating from those domains violate policies in a manner that would cause his system to begin rejecting e-mail traffic until such time as the noncompliant domain returns to compliance. Thus, a detailed review of Srivastava's disclosure that the Examiner relies on for rejection reveals that it is entirely silent regarding any feature that could function in a manner that is analogous to Applicant's flow control mechanism. In order to achieve Applicant's result (e.g., blocking spam and denial of service attacks) in Srivastava's system, one would have to add Applicant's filtering mechanism to Srivastava's system.

Further, the Examiner points to Srivastava's Col. 5, lines 9-15, which states:

In the described embodiment, access to the message store 304 is multithreaded thereby allowing a single process to manage a large number of connections since each connection is handled by a thread. In this way, multithreaded access maximizes both performance and scalability by minimizing the system resources required for the management of each connection.

Here, the Examiner contends that Srivastava's multithreaded access to a single message store is the same as Applicant's approach of spawning a second process for handling a

new incoming connection.

"Threads" and "processes" are not the same. A "process" is an executing program or task. A "thread" is a part of a process that can execute independently of other parts; it exists within a process and uses the process' resources. Unlike processes, multiple threads run within the same address space and share their process' data. The concepts of threads and processes are well known and well documented in the technical literature. See, e.g., Kalev, Danny, "Processes and Threads," ITWorld.com, February 9, 2001, a copy of which is attached for the Examiner's convenience. The article discusses threads and processes in the context of the Linux operating system, but the discussed concepts apply equally well to other operating systems (e.g., UNIX, Windows, Macintosh OS X). The article is not intended to be an Information Disclosure Statement, but is instead general material merely offered to assist the Examiner in understanding the differences between the well-known concepts of threads and processes.

Without discussing the other deficiencies of Srivastava at this point (e.g., a single "message store" in Srivastava versus multiple incoming connections from different domains in Applicant's system), it is clear that the section that the Examiner cites in Srivastava discusses the use of multiple threads, not the spawning of additional processes. If anything, Srivastava at this point teaches away from Applicant's claimed approach of multiple processes (not Srivastava's approach of a single process with multiple threads).

Applicant's flow control invention provides a facility for moderating the flow of SMTP traffic (connections, aggregate volume, and unique senders) into a server or set of servers. This feature is brought out in Applicant's claims. For example, claim 1 recites:

receiving at a first process a request from a particular domain to establish a new connection for transmitting a particular e-mail message to the e-mail system;

(Emphasis added.)

This is an incoming connection from another domain (particularly, an MTA at another domain), for the purpose of doing an MTA-to-MTA email exchange. This is not an operation for servicing user requests.

Further, claim 1 requires:

in response to receipt of said request from the particular domain, creating a second process for handling the request to establish a new connection, said second process being connected to a flow control filter providing filtering on a per-domain basis;

comparing the request from the particular domain against configurable policy rules; and

denying the request if any of said policy rules would be violated.

(Emphasis added.)

Here, the claim requires that the above-mentioned incoming connection is passed through a domain-specific filter. This approach allows Applicant's flow control invention to detect and prevent massive spam from being received on incoming connections of a particular domain. Srivastava's approach of granting user services cannot be morphed into system that prevents incoming massive spam from an MTA of a particular domain.

Srivastava is essentially a method for providing virtual hosting services for e-mail and web pages, with the ability to create virtual users within that context and optionally delegate authority to those users to manage parts of the virtual space so provisioned. On startup, Applicant's flow control invention reads a configuration file and then reacts to SMTP traffic it observes. It has no "user-serviceable parts"; only the e-mail administrator has access to read or change its configuration. Although the two approaches converge insofar as they are both related to providing e-mail service at ISPs and can do some amount of per-user validity checking, the convergence ends there. They otherwise operate on different types of data and in different layers of the OSI protocol stack. It is respectfully submitted that these distinctions are apparent from Applicant's claims and that the claimed invention distinguishes over Srivastava.

The other independent claims rejected under Section 102 (i.e., claims 21 and 41) include the above-mentioned distinguishing per-domain filtering or policy enforcement claim limitations, and are therefore believed to be allowable for the reasons stated above. (The dependent claims rejected under Section 102 are believed to be allowable by virtue

of depending from the foregoing independent claims.) Accordingly, it is respectfully submitted that the claims distinguish over Srivastava, and set forth a patentable invention under Section 102.

B. Section 103: Srivastava and Spam!

Claims 6, 12, 14, 29-30 and 46 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! (Cranor and LaMacchia, Communications of the ACM, August 1998). Here, the Examiner relies on Srivastava as above, and adds Spam! for the proposition that it is obvious to add a spam filter.

The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava in its failure to teach or suggest Applicant's per-domain flow control filter. Adding Spam! does not cure these deficiencies. Further, the claims are believed to be allowable for the following additional reasons.

At the outset, it is important to recognize that Applicant does not claim to be the first to invent a spam filter, and Applicant's invention itself is not a spam filter but instead a flow control filter which operates at server level (MTA) to monitor the the behavior of different domains that are connecting to send incoming e-mail. For example, a "bad" behavior would include a denial-of-service attack, which of course itself is not "spam" (unsolicited e-mail) as that term is generally understood.

With respect to claim 6, for example, Applicant's claim discusses comparing the sender information against policy rules for a particular domain. This is not simply rejecting an e-mail piece based on it coming from a sender that has been blacklisted. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following teaching from Applicant's specification (at page 17, line 20 to page 18, line 5):

As described above, with each new connection a child MTA process is created. In accordance with the present invention, each child process connects to the flow control filter service, so that it can interact with the service during arrival of a message. This interaction provides a complete description of the incoming client, including IP address and host name, as well as the complete SMTP interaction, including HELO (i.e., initial "hello"

handshake), MAIL FROM (i.e., sender information), RCPT TO (i.e., recipient list), and DATA (i.e., entire message body). Since the flow control filter service monitors all children processes, it attains a global view of traffic flowing through the system. **By virtue of its global view, the flow control filter service can track information on a per domain basis, including total volume of e-mail received from a particular domain over a given period of time. Examples of other metrics that may be tracked include total connections and total senders (count) encountered for a particular domain over a given period of time. Other examples include total number of different recipients, total number of envelopes, and total aggregate volume of mail encountered for a particular domain over a given period of time.** Since the knowledge lost by the forking process is captured by the flow control filter service, the service is in a position to enforce policy-based rules, including placing restrictions on child processes, based on the per-domain tallies encountered.

(Emphasis added.)

As described above, the purpose of examining header information is not to accept or reject a given single piece of e-mail based on spam criteria (e.g., blacklisted or not), but is use in conjunction with Applicant's flow control filter to further characterize the given domain that is being monitored. For example, as specified in the passage above, one of the criteria may be "number of different recipients." This requires the system to look at e-mail header information, but note in particular that the header information is not being examined for purposes of identifying an individual piece of e-mail as spam, but instead is being used to further characterize the current behavior of the given domain that is being monitored (in order to determine whether server-level intervention is warranted).

With respect to claim 12, Applicant's claim discusses comparing the e-mail message body data against policy rules for a particular domain. This is not simply rejecting an e-mail piece based on it having certain content (e.g., explicit content) that is detected and rejected by a spam filter. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Consider the following from Applicant's specification (at page 10, line 27 to page 11, line 2):

The MTA reports the message body, which may be transmitted as one or more blocks. The method updates a running total of message size. This information is used to determine the aggregate total of bytes received from a given source over a period of time. The MTA reports end of message for the current incoming message. The method compares the message size against class limits specified in the configuration file. Again as before, if specified limits are exceeded, the method terminates with the filter rejecting the message (returning any administrator-defined error code).

(Emphasis added.)

As described above, the message header may be examined, not for the purpose out of accepting or rejecting a given single piece of e-mail based on spam criteria (e.g., offensive content), but is used in conjunction with Applicant's flow control filter to further characterize the given source -- a particular domain -- that is being monitored. For example, as specified in the passage above, one of the criteria may be "aggregate total bytes received from a given source over a period of time." This requires the system to look at the message body, but note in particular that the message body is not being examined for purposes of identifying the e-mail as having spam content, but instead is being used to further characterize the current behavior of the given domain that is being monitored.

Regarding claim 14, for example, the Examiner contends that this is taught by Spam! at page 79, which indicates that ISPs may limit the number of outbound messages each subscriber can send. However, that is not Applicant's claim limitation. Instead, claim 14 recites (in amended form) that "said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a desired period of time." As readily apparent from the claim language, the volume of e-mail being regulated is that coming from a given domain, not that coming from an individual user or subscriber. Thus, for example, applying Applicant's invention, the *uspto.gov* e-mail server could be configured to detect an abnormal amount of e-mail coming from *aol.com* (and take appropriate action, accordingly), for example as a result of a denial-of-service attack originating from that domain. Such a result cannot be achieved by simply using a

spam filter at the ISP (*aol.com*) for attempting to detect an abnormally high level of e-mail from any given subscriber. And, in the case of a distributed denial-of-service attack, the attack may be distributed over numerous subscriber accounts (e.g., as a result of Trojan/zombie infection), and it may very well be the case that no one subscriber has an abnormal level of outbound messages.

Regarding claims 29 and 30, for example, the Examiner contends that Spam! at p. 79 teaches that limits can be placed on a domain. However, Spam! itself describes placing limits on individual subscriber accounts. Spam! only describes blocking e-mail from a bogus domain. It contains no description of how a spam filter could continually monitor the e-mail traffic behavior of a given domain, and apply configurable policy rules. Again, the characteristics of e-mail traffic coming from a given domain (e.g., *aol.com*) are not the same as the characteristics of e-mail traffic coming from an individual subscriber (e.g., *john.smith@aol.com*).

To establish a prima facie case of obviousness under Section 103, the Examiner must establish, among other things, that the prior art reference (or references when combined) must teach or suggest all the claim limitations. (See e.g., MPEP 2142). For the reasons stated above, the references cited by the Examiner fail to meet these conditions. Accordingly, it is respectfully submitted that the claims distinguish over the references and are patentable under Section 103.

C. Section 103: Srivastava and RFC 821

Claim 35 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and RFC 821. The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava in its failure to teach or suggest Applicant's per-domain flow control filter. Adding RFC 821 does not cure these deficiencies. Further, the claim is believed to be allowable for the following additional reasons.

As noted above for claim 6, Applicant's invention uses the sender information to monitor e-mail traffic behavior from a given domain (and take action based on configurable policy rules applicable for that given domain). This is not simply processing

an individual e-mail piece based on a particular sender that it is coming from. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain. Accordingly, the claim is believed to distinguish over the art, and therefore be allowable under Section 103.

D. Section 103: Srivastava, Spam!, and RFC 821

Claims 7, 13, 36, and 39-40 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of RFC 821. The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava and Spam! The addition of RFC 821 does not remedy these deficiencies.

Regarding claim 36, it should be noted that the claim limitation recites that "said flow control filter examines said sender information to ascertain whether any of said rules would be violated." Base claim 21 provides "a flow control filter, in communication with said child processes and said set of rules, providing filtering based on each domain's conformance to said rules." (Emphasis added.) Again, the plain language of the claim indicates that the sender information is being used to characterize the e-mail traffic behavior of a given domain. This is not the same functionality as an ISP or end-user spam filter blocking e-mail based on a particular undesired sender. Similarly, regarding claim 40, it is apparent that the claimed invention is using the e-mail body to characterize the e-mail traffic behavior of a given domain -- a functionality that is not present in ISP or end-user spam filters.

E. Section 103: Srivastava and Apache

Claim 26 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Apache HTTP Server Configuration Files ("Apache"). The claim is believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava, which are not cured by the addition of Apache.

F. Section 103: Srivastava and Mosberger

Claims 2 and 42 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Mosberger et al. (U.S. Patent 6,438,597, "Mosberger"). The claim is believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava, which are not cured by the addition of Mosberger. Mosberger describes firewall-like features of controlling connections (e.g., based on domain). However, Mosberger does not provide sufficient teaching to morph Srivastava's virtual domain hosting into an e-mail flow control filter that controls connections based on domain-specific behavior of e-mail traffic.

G. Section 103: Srivastava and Shaw

Claim 9 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw et al. (U.S. Patent No. 6,282,565, "Shaw"). The claim is believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava, which are not cured by the addition of Shaw. Applicant's invention uses the recipient information to monitor e-mail traffic behavior from a given domain (and take action based on configurable policy rules applicable for that given domain). This is not simply processing an individual e-mail piece based on a particular recipient that it is going to. Instead, this is used in the context of other criteria that have been established in the configurable policy rules for that particular domain.

H. Section 103: Srivastava, Spam!, and Shaw

Claim 15 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 14 above, and further in view of Shaw. The claim is believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava and Spam!, which have been extensively described above. These deficiencies are not remedied by Shaw. Further, the Examiner relies on Shaw for disclosing the claim limitation of "limiting the size of incoming e-mail messages based on a maximum number of bytes." However, this is not what Applicant's claim states. Instead, claim 15 recites: "said **maximum aggregate volume** is based on total byte count of e-mail received from a given domain over a desired period of time." (Emphasis

added.) Claim 14 (claim 15's parent) recites: "wherein said configurable policy rules specify a maximum aggregate volume of e-mail permitted by a given domain over a desired period of time." Limiting the maximum average volume from a given domain is not the same as limiting the size of a given incoming e-mail message. Accordingly, the cited art does not serve as an appropriate rejection.

I. Section 103: Srivastava, Shaw, and Sash

Claims 11 and 44 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of Sash (U.S. Pub. No. 2003/0167250). The claim is believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava and Shaw, which are not cured by Sash. Here, the Examiner adds Sash for the additional teaching of "a maximum number of different recipients permitted..." Sash describes an information template and describes limiting a maximum number of recipients that an information template can be sent to (i.e., limit the number of times it can be forwarded to other recipients). Applicant's claim limitation states, "said configurable policy rules specify a maximum number of different recipients permitted by a given domain over a desired period of time." This would apply, for instance, in this scenario of e-mail traffic coming from a particular domain (e.g., *advertiser.net*) having an inordinate number of different recipients (say, e.g., > 10 million). Placing a restriction on the number of times that a data object can be forwarded (e.g., Sash's restriction on the number of recipients that Sash's information template can be forwarded to) bears little relevance to Applicant's claim limitation.

J. Section 103: Srivastava, Spam!, and RFC 821

Claim 10 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 9 above, and further in view of RFC 821. The claim is believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava and Spam!, which are not cured by RFC 821.

K. Section 103: Srivastava and Rakoshitz

Claims 4 and 5 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Rakoshitz et al. (U.S. Patent No. 6,816,903, "Rakoshitz"). The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava which Rakoshitz does not cure. The claims are believed to be allowable for the following additional reasons.

Rakoshitz describes "select counters for monitoring incoming and outgoing traffic from a link" (Rakoshitz, at Col. 21, lines 2-3). Nowhere does Rakoshitz describe maintaining "a counter indicating how many connections have been granted to the particular domain" (emphasis added), as required by Applicant's claim. To the extent that Rakoshitz teaches a counter, the Rakoshitz counter is one that tracks individual links. In particular, no description is given which teaches or suggests that the individual links traceable to or referencing a particular domain be tracked with a counter. Accordingly, at best, Rakoshitz teaches away from Applicant's domain counter claim limitation.

L. Section 103: Srivastava, Spam!, and Bates

Claims 8, 43, 45, and 49 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Spam! as applied to claim 6 above, and further in view of Bates et al. (U.S. Patent No. 6,779,021, "Bates"). The claims are believed to be allowable for at least the reasons cited above pertaining to the deficiencies of Srivastava and Spam!, which Bates does not cure. The claims are believed to be allowable for the following additional reasons.

The passage cited by the Examiner comes from the Background Section of Bates. There, Bates describes basic spam filtering techniques, such as blocking on sender. Applicant's claim limitation, however, requires: "a maximum number of different senders permitted by a given domain over a desired period of time" (emphasis added). A review of Bates indicates no such feature described or suggested. Further, to the extent that Bates repeats spam filtering information about blocking without regard to the behavior of a particular domain, Bates teaches away from Applicant's claimed approach.

M. Section 103: Srivastava, Shaw, and RFC 821

Claim 10, 37, and 38 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Srivastava and Shaw as applied to claim 9 above, and further in view of RFC 821. The claims are believed to be allowable for lease these reasons cited about pertaining to Srivastava and Shaw -- that is, Applicant's use of the recipient information to monitor e-mail traffic behavior from a given domain. RFC 821 does not remedy this deficiency of the combined references.

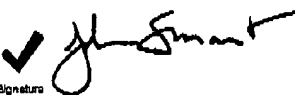
Conclusion

In view of the foregoing remarks and the amendment to the claims, it is believed that all claims are now in condition for allowance. Hence, it is respectfully requested that the application be passed to issue at an early date.

If for any reason the Examiner feels that a telephone conference would in any way expedite prosecution of the subject application, the Examiner is invited to telephone the undersigned at 408 884 1507.

Respectfully submitted,

Date: April 1, 2005


Digitally signed by John A. Smart
Date: 2005.04.01 14:57:28 -0800
✓
Signature Valid

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX